

The Role of Innovation within Economics

Russell K. Standish

High Performance Computing Support Unit
The University of New South Wales,
Sydney, 2052

February 8, 2008

1 Introduction

Comparative statics is a special case of dynamics, in which a unique stable equilibrium is assumed to exist, and two equilibria — determined by different parameter values — are compared. Similarly, dynamics is a special case of an evolutionary process, in which the degrees of freedom of the system are held constant. Marshall, whose *Principles*[16] played a large rôle in setting economics upon the comparative statics course, was nonetheless aware of the primacy of the evolutionary paradigm when he wrote on page xiv that:

The Mecca of the economist lies in economic biology rather than in economic dynamics. But biological conceptions are more complex than those of dynamics; a volume on Foundations must therefore give a relatively large place to mechanical analogies; and frequent use of the term ‘equilibrium’, which suggest something of a statical analogy.... The modern mathematician is familiar with the notion that dynamics includes statics. If he can solve a problem dynamically, he seldom cares to solve it statically also... But the statical solution has claims of its own. It is simpler than the dynamical; it may afford useful preparation and training for the more difficult dynamical solution; and it may be the first step towards a provisional and partial solution in problems so complex that a complete dynamical solution is beyond our attainment.

One century later, intellectual progress has enabled us to contemplate the construction of evolutionary models of processes which could once only be modeled statically.

Early ecological theories followed a similar path of development. Interactions between organisms were studied in order to determine the population densities

that kept the ecology in equilibrium and the only form of dynamics considered was instantaneous movement from one equilibrium to another. It was always assumed that an ecology was always near a stable equilibrium, in spite of counter examples such as the stable limit cycles in the Lotka-Volterra predator-prey system[19].

In the 1960s, interest in dynamical systems was reawakened, and the study of *Chaos* was born. These techniques were applied to ecology in the 1970s, and it was realised that the situation was vastly more complicated, with not only limit cycles being possible, but also entities called *strange attractors*. Over the years, appreciation has grown that the stable equilibrium is more of an exception than the norm.

In the mid 1980s, Raup and Sepkowski[24] noticed certain statistical regularities with the pattern of speciation and extinction in the fossil record. These were later shown to be in the form of a power law[?, ?], i.e the pattern of extinctions was such that the frequency of an extinction event of size x is proportional to $x^{-\alpha}$, where α is a small positive constant. Power laws crop up in many different areas.[25]

Per Bak[3] studied a model of a sandpile that had a continuous stream of sand added to the top. Bak's process goes by the name of *self-organised criticality*, and he, along with Kauffman[13] promote this as an explanation for Raup's data. It now seems likely that evolution is an endogenously self-organised critical process¹

In my work, which I will introduce later, I have demonstrated the existence of self-organised criticality in a model evolutionary ecology. The criticality is quite a robust feature over a wide range of input parameters. This indicates that that evolutionary systems are typically endogenously critical.

In summary, we have a hierarchy of approaches, from the static, to the dynamic to the evolutionary:

$$\text{statics} \subset \text{dynamics} \subset \text{evolution}$$

The state of ecological thought has followed this chain from left to right, as computational techniques have improved to embrace the computationally more difficult dynamical models, and then the even more computationally difficult evolutionary models. It is to be hoped that the same path will be followed by economic modeling.

There is a clear parallel between the development of ecological thought and that of economic thought. Dynamics was introduced to economics and championed by Kaldor, Goodwin et al. in the 1950s and 60s; and boosted by the developments in nonlinear analysis in the 1980s, through the work of people

¹Alternatively, perhaps exogenous influences (such as volcanos, meteorites, climatic fluctuations, supernovae etc) obey a power law spectrum. It is quite possible that both endogenous and exogenous effects contribute to the statistical properties of biological evolution and that a future research program will look at untangling the two effects. To this effect, Newman[?] provides an interesting perspective.

like Blatt[5]. This volume testifies to a growing industry of dynamical economic modeling. However, perhaps now is the time to embark on the next rung up the ladder, and approach economics from an evolutionary point of view. Then perhaps we might be catching a glimpse of the Mecca that Marshall referred to a hundred years ago.

This paper first considers a general dynamical system undergoing evolution of the determining equation, then outlines *Ecolab*, a model of species interaction undergoing evolution, finally introducing a possible economics model based on the von Neumann model with evolution of the interaction coefficients.

2 Linearisation of a Dynamical System

Our launchpad for a theory of evolutionary systems is dynamical systems theory[8]. Typically this will be manifested in a first order nonlinear differential equation of the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$$

where $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{f} : \mathbb{R} \rightarrow \mathbb{R}^n$. The dot refers to the derivative with respect to time.

Dynamical systems theory starts by considering the equilibria of the system, i.e. the points $\hat{\mathbf{x}}$ such that $\mathbf{f}(\hat{\mathbf{x}}) = 0$. Then in the neighbourhood of $\hat{\mathbf{x}}$, the behaviour of the system is determined by the linear approximation

$$\dot{\mathbf{x}} = D\mathbf{f}|_{\hat{\mathbf{x}}} \cdot (\mathbf{x} - \hat{\mathbf{x}}).$$

The stability of $\hat{\mathbf{x}}$ is determined by the negative definiteness of $D\mathbf{f}|_{\hat{\mathbf{x}}}$ ². This condition imposes n inequalities on the system constraining the form of \mathbf{f} . There may additionally be a further n inequalities for $\hat{\mathbf{x}}$ to be a meaningful solution, for example if the components of $\hat{\mathbf{x}}$ are production values, every component of $\hat{\mathbf{x}}$ must be non-negative.

General equilibrium economics has attempted to find the conditions under which a unique, stable equilibrium will exist. This needn't be the case, and interesting (i.e. bounded) behaviour can take place around unstable equilibria, in the form of limit cycles or even the strange attractors beloved of chaos theorists. A favourite model of the latter researchers is the logistic equation, which first arose in a biological context[17], but has been applied to economics[31] amongst other things.

That limit cycles and chaotic behaviour can be observed in economics is a view that should have by now been accepted[5]. However, the question remains as to whether this behaviour is pathological, i.e. whether linear neoclassical theory is applicable in most cases, and the only remaining difficulty is determining if linear theory applies to a specific economy (the “econometric problem”), or whether chaotic behaviour is indeed the norm.

² $D\mathbf{f}|_{\hat{\mathbf{x}}}$ is negative definite if $\mathbf{x} \cdot D\mathbf{f}|_{\hat{\mathbf{x}}} \cdot \mathbf{x} < 0$ for all \mathbf{x} . This also implies that all eigenvalues of $D\mathbf{f}$ have negative real part.

3 Limits of Linear Economic Theory

Returning to the question of stability, the fact that $2n$ inequalities must be satisfied would imply that a randomly chosen n -dimensional economics model would have a probability of 4^{-n} of a given equilibrium being stable. The situation does look bleak, but economics are not generated randomly in the real world, rather they are the result of an evolutionary process. We need to examine the process of cultural evolution to answer this question.

The analogue of mutation in biology would be innovation in economics, as a new process or technique introduced into production, or as a new form of marketing, or a new company with a somewhat unusual approach to doing business. The effect of innovation is to add new degrees of freedom to the dynamical system, which usually will destabilise the system. The system will then tend to evolve so as to lose some of the degrees of freedom, by for example old production techniques being abandoned, or companies going bankrupt. This is analogous to species becoming extinct in the natural world.

Lets consider what happens to the largest eigenvalue of $Df|_{\hat{x}}$. Suppose initially, the system has a stable equilibrium, in which case all the eigenvalues have negative real part. As innovations are added to the system, the largest eigenvalue will increase towards zero. As it passes zero, the system destabilises, and the system will start to exhibit limit cycles or chaotic behaviour. As further innovations are added to the system, a property called *permanence*³ is no longer satisfied, and some event such as a bankruptcy will occur to remove active processes from the system. This will restore permanency to the system, and possibly even stability. Such a process is called *self-organised criticality*[3] which gives rise to a power law spectrum of the booms and busts, successful innovations and bankruptcies.⁴

4 Ecolab and the Dynamics of Evolution

This section outlines a model of an evolving ecology[29, 27] that is analogous to an economic system with input-output relations of production and product innovation. The ecology is described by a generalised Lotka-Volterra equation, which is perhaps the simplest ecological model to use.

$$\dot{n}_i = r_i n_i + \sum_{j=1}^{n_{sp}} \beta_{ij} n_i n_j \quad (1)$$

³that there is a set of points \mathbf{x}_0 whose trajectories $\mathbf{x}(t)$ always remain away from the boundary, i.e. $x_i(t) > \delta \ \exists \delta > 0$

⁴This really implies that conclusion derived from comparative static economic analysis are almost never valid, except perhaps on sufficiently small time scales while the maximum eigenvalue of $Df|_{\hat{x}}$ is negative.

Here \mathbf{r} is the difference between the birth rate and death rate for each species, in the absence of competition or symbiosis. $\boldsymbol{\beta}$ is the interaction term between species, with the diagonal terms referring to the species' self limitation, which is related in a simple way to the carrying capacity K_i for that species in the environment by $K_i = -r_i\beta_{ii}$. In the literature (eg Strobeck[30], Case[7]) the interaction terms are expressed in a normalised form, $\alpha_{ij} = -K_i/r_i\beta_{ij}$, and $\alpha_{ii} = 1$ by definition. \mathbf{n} is the species density.

These equations are simulated on a simulator called *Ecolab*.[28] The vectors \mathbf{n} and \mathbf{r} are stored as dynamic arrays, the size of which (i.e. the system dimension) can change in time. The interaction array is stored in row/column sparse form, consisting of the four arrays `diag`, `val`, `row` and `col`. Equation (1) can be written as:

```
tmp[row] = beta.val * n[beta.col];
n += (r + beta.diag + tmp) * n;
```

This code makes up the `generate` operator in the Ecolab system. Other operators include `compact`, which removes species that have become extinct from the system (to optimise computational performance) and `mutate`, which adds a certain number of new species to the system, according to a specific algorithm to be discussed later. The operators can be called from a scripting language called `TCL`[21], that allows different types of experiments to be performed without recompiling the code.

Before discussing the mutation algorithm in more detail, equation (1) must be analysed to determine the conditions $\boldsymbol{\beta}$ must satisfy for the system to be real, and also to determine the different regimes of dynamics, from the linear (stable equilibrium) case, to limit cycles and chaos to the actual breakdown of the ecosystem.

4.1 Linear Analysis

Linear analysis starts with the fixed point of equation (1)

$$\hat{\mathbf{n}} = -\boldsymbol{\beta}^{-1}\mathbf{r}, \quad (2)$$

where $\hat{\mathbf{n}} = 0$. There is precisely one fixed point in the interior of the space of population densities (i.e. \mathbf{n} such that $n_i > 0$) provided that all components of $\hat{\mathbf{n}}$ are positive, giving rise to the following inequalities:

$$\hat{n}_i = (\boldsymbol{\beta}^{-1}\mathbf{r})_i > 0, \quad \forall i \quad (3)$$

This interior space is denoted $\mathbb{R}_+^{n_{sp}}$ mathematically.

There may also be fixed points on the boundary of $\mathbb{R}_+^{n_{sp}}$, where one or more components of \mathbf{n} are zero (corresponding to an extinct species). This is

because the subecology with the living species only (i.e. with the extinct species removed) is equivalent to the full system.

The *stability* of this point is related to the negative definiteness of derivative of $\dot{\mathbf{n}}$ at $\hat{\mathbf{n}}$. The components of the derivative are given by

$$\frac{\partial \dot{n}_i}{\partial n_j} = \delta_{ij} \left(r_i + \sum_k \beta_{ik} n_k \right) + \beta_{ij} n_i \quad (4)$$

Substituting eq (2) gives

$$\frac{\partial \dot{n}_i}{\partial n_j} \Big|_{\hat{\mathbf{n}}} = -\beta_{ij} (\boldsymbol{\beta}^{-1} \mathbf{r})_i \quad (5)$$

Stability of the fixed point requires that this matrix should be negative definite. Since the $(\boldsymbol{\beta}^{-1} \mathbf{r})_i$ are all negative by virtue of (3), this is equivalent to $\boldsymbol{\beta}$ being negative definite, or equivalently, that its n_{sp} eigenvalues all have negative real part. Taken together with the inequalities (3), this implies that $2n_{sp}$ inequalities must be satisfied for the fixed point to be stable. This point was made by Strobeck[30], in a slightly different form. (Note that Strobeck implicitly assumes that $\sum_i r_i \hat{n}_i / K_i > 0$, so comes to the conclusion that $2n_{sp}-1$ conditions are required.) If one were to randomly pick coefficients for a Lotka-Volterra system, then it has a probability of $4^{-n_{sp}}$ of being stable, i.e. one expects ecosystems to become more unstable as the number of species increases[18].

4.2 Permanence

Whilst stability is a nice mathematical property, it has rather less relevance when it comes to real ecologies. For example the traditional predator-prey system studied by Lotka and Volterra has a limit cycle. The fixed point is decidedly unstable, yet the ecology is *permanent* in the sense that both species' densities are larger than some threshold value for all time. Hofbauer et al. [9] and Law and Blackford[15] discuss the concept of *permanence* in Lotka-Volterra systems, which is the property that there is a compact absorbing set $\mathcal{M} \subset \mathbb{R}_+^{n_{sp}}$ i.e once a trajectory of the system has entered \mathcal{M} , it remains in \mathcal{M} . They derive a sufficient condition for permanence due to Jansen[11] of the form:

$$\sum_i p_i f_i(\hat{\mathbf{n}}_B) = \sum_i p_i (r_i - \sum_j \beta_{ij} \hat{n}_{Bj}) > 0, \quad \exists p_i > 0 \quad (6)$$

for every $\hat{\mathbf{n}}_B$ equilibrium points lying on the boundary ($\hat{n}_{Bi} = 0 \quad \exists i$), provided the system is *bounded* (or equivalently *dissipative*). This condition is more general than stability of the equilibrium — the latter condition implies that a local neighbourhood of the equilibrium is an absorbing set. Also, the averaging property of Lotka-Volterra systems implies that the equilibrium must lie in the positive cone $\mathbb{R}_+^{n_{sp}}$. So (3) must still hold for permanence.

Consider the boundary points $\hat{\mathbf{n}}_B$ that are missing a single species i . Then Jansen's condition for these boundary points is

$$r_i - \sum_j \beta_{ij} \hat{n}_{Bj} > 0. \quad (7)$$

This set of conditions is linearly independent. Let the number of such boundary points be denoted by $n_B \leq n_{sp}$. Then the set of conditions (6) will have rank $n_B \leq \nu \leq n_{sp}$ (the number of linearly independent conditions), so the system has at most probability $2^{-n_{sp}-\nu}$ of satisfy Jansen's permanence condition if the coefficients are chosen uniformly at random. As stability is also sufficient for permanence, the probability lies between $4^{-n_{sp}}$ and $2^{-n_{sp}-\nu}$.

Another rather important property is *resistance to invasion*.[7] Consider a boundary equilibrium $\hat{\mathbf{n}}_B$. If it is proof against invasion from the missing species, then the full system cannot be permanent. For the boundary points that miss a single species, this implies that condition (7) is necessarily satisfied for permanence, along with (3). The probability of permanence is then bounded above by $2^{-n_{sp}-n_B}$.

Thus whilst a randomly selected ecology is more likely to be permanent than to have a stable equilibrium, the likelihood decreases exponentially with increase in species number.

4.3 Boundedness

It is necessary that the ecology be *bounded*, ie that $\sum n_i < N \ \exists N, \ \forall t > 0$. This requires

$$\sum_i \dot{n}_i = \mathbf{r} \cdot \mathbf{n} + \mathbf{n} \cdot \boldsymbol{\beta} \mathbf{n} < 0, \ \forall \mathbf{n} : \sum_i n_i > N \ \exists N \quad (8)$$

As \mathbf{n} becomes large in any direction, this functional is dominated by the quadratic term, so this implies that

$$\mathbf{n} \cdot \boldsymbol{\beta} \mathbf{n} \leq 0 \ \forall \mathbf{n} : n_i > 0. \quad (9)$$

If strict equality holds, then $\mathbf{r} \cdot \mathbf{n} < 0$. Negative definiteness of $\boldsymbol{\beta}$ is sufficient, but not necessary for this condition. Another sufficient condition is to require $\forall i, j, \beta_{ii} < 0$ and $\beta_{ij} + \beta_{ji} \leq 0$, which is used in the current study. This condition is satisfied by the Predator-Prey equations, and so does allow multi-trophic systems to be built, but does not allow the possibility of symbiosis. Its main advantage is its simplicity of implementation, along with the range of interesting (i.e. non limit point) behaviour it encompasses.

4.4 Mutation

Adding mutation involves adding an additional operator to equation (1)

$$\dot{\mathbf{n}} = \mathbf{r} * \mathbf{n} + \mathbf{n} * \boldsymbol{\beta} \mathbf{n} + \text{mutate}(\boldsymbol{\mu}, \mathbf{r}, \mathbf{n}) \quad (10)$$

where $*$ refers to elementwise multiplication.

The mutation operator must generate new degrees of freedom $i > n_{sp}$ (where n_{sp} is the number species currently in the ecology), somehow defining the new ecological coefficients $\{r_i | i > n_{sp}\}, \{\beta_{ij} | i > n_{sp} \text{ or } j > n_{sp}\}$ from the previous state of the system. In reality, there is another layer (hidden in equation (1) called the genotypic layer, where each organism has a definite genotype. There is a specific map from the genotypic layer to the space of ecological coefficients (hereafter called the phenotypic layer) called the *embryology*. Then the mutation operator is a convolution of the genetic algorithm operations operating at the genotypic layer, with the embryology.

A few studies, including Ray's Tierra world, do this with an explicit mapping from the genotype to some particular organism property (e.g. interpreted as machine language instructions, or as weight in a neural net). These organisms then interact with one another to determine the population dynamics. In this model, however, we are doing away with the organismal layer, and so an explicit embryology is impossible. The only possibility left is to use a statistical model of embryology. The mapping between genotype space and the population parameters r, β is expected to look like a rugged landscape, however, if two genotypes are close together (in a Hamming sense) then one might expect that the phenotypes are likely to be similar, as would the population parameters. This I call *random embryology with locality*. Here, we tend to idealise genotypes as bit strings, although strings over an arbitrary alphabet (eg the four DNA bases ACGT) can equally be considered.⁵

In the simple case of point mutations, the probability $P(x)$ of any child lying distance x in genotype space from its parent follows a Poisson distribution, as this is the distribution of the number of bit flips, or deletions that might occur with a point mutation. Random embryology with locality implies that the phenotypic parameters are distributed randomly about the parent species, with a standard deviation that depends monotonically on the genotypic displacement. The simplest such model is to distribute the phenotypic parameters in a Gaussian fashion about the parent's values, with standard deviation proportional to the genotypic displacement. This constant of proportionality can be conflated with the species' intrinsic mutation rate, to give rise another phenotypic parameter μ . It is assumed that the probability of a mutation generating a previously existing species is negligible, and can be ignored. We also need another arbitrary parameter ρ , "species radius", which can be understood as the minimum genotypic distance separating species, conflated with the same constant of proportionality as μ .

We may represent the Ecolab embryology as a probability distribution $f(p, g) = \sqrt{\frac{2}{\pi}} \frac{\mu e^{-(\frac{g}{2\rho})^2}}{g}$, where $p = |r_i - r_j|/|r_i|$ or $p = |\beta_{ik} - \beta_{jk}|/|\beta_{ik}|$ is the distance between two species' phenotypic parameters, and g is the difference between the

⁵The Hamming distance is the number of bits (bases) that differ between the two strings. So for example if a single bit has been removed from one string, the Hamming distance is one.

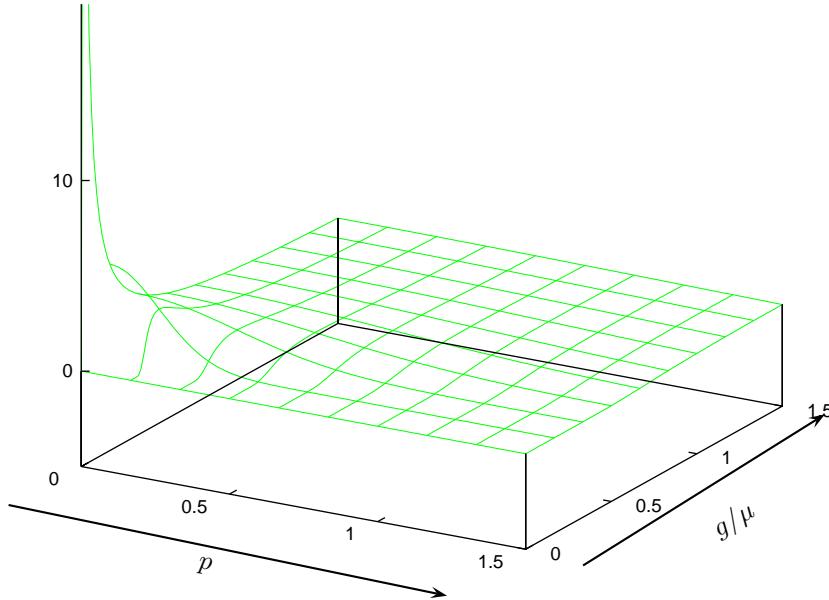


Figure 1: Probability distribution of the relation between genotype difference and the corresponding phenotype difference

two genotypes. Figure 1 shows the general form of this probability distribution.

Figure 2 shows the probability distribution of a mutant phenotypical coefficient about that of its parent's value. This is given by

$$\int_0^\infty \sqrt{\frac{2}{\pi}} \frac{e^{-g/\mu - (\frac{\mu p}{2g})^2}}{g} \mu dg. \quad (11)$$

In summary, the mutation algorithm is as follows:

1. The number of mutant species arising from species i within a timestep is $\mu_i r_i n_i / \rho$. This number is rounded stochastically to the nearest integer, e.g. 0.25 is rounded up to 1 25% of the time and down to 0 75% of the time.
2. Roll a random number from a Poisson distribution $e^{-x/\mu+\rho}$ to determine the standard deviation σ of phenotypic variation.
3. Vary \mathbf{r} according to a Gaussian distribution about the parents' values, with σr_0 as the standard deviation, where r_0 is the range of values that \mathbf{r} is initialised to, ie $r_0 = \max_i r_i|_{t=0} - \min_i r_i|_{t=0}$.

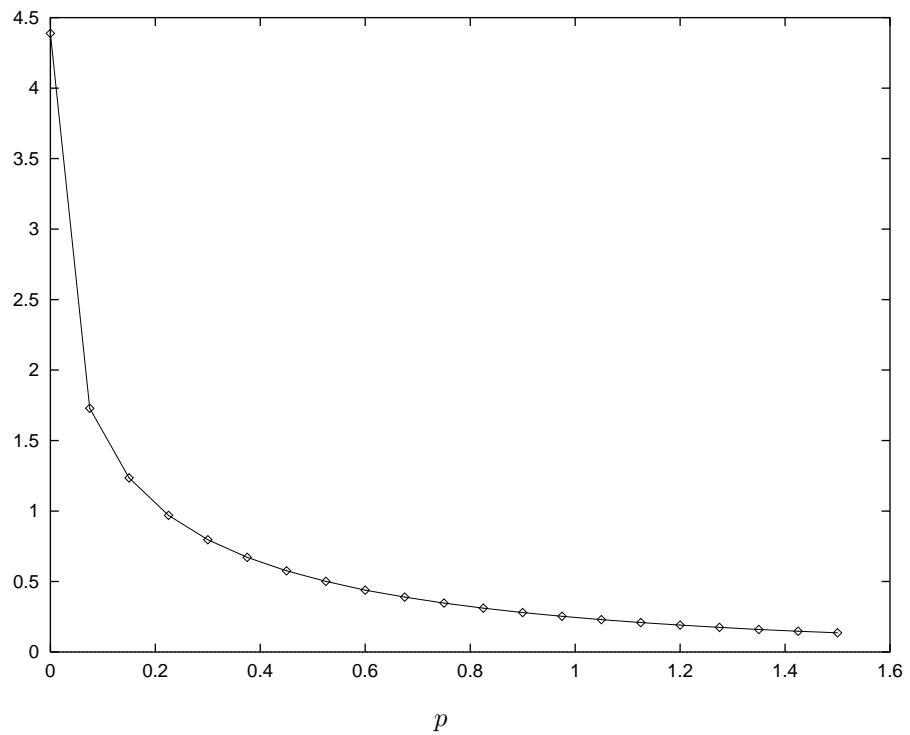


Figure 2: The probability distribution of a mutant phenotypical coefficient about that of its parent's value. This has been calculated by numerical integration from formula (11). Note that the curve actually diverges at 0.

4. The diagonal part of β must be negative, so vary β according to a log-normal distribution. This means that if the old value is β , the new value becomes $\beta' = -\exp(-\ln(\beta) + \sigma)$. These values cannot arbitrarily approach 0, however, as this would imply that some species make arbitrarily small demands on the environment, and will become infinite in number. In ecolab, the diagonal interactions terms prevented from becoming larger than $-r/(.1 * \text{INT_MAX})$, where r is the corresponding growth rate for the new species.
5. The off diagonal components of β , are varied in a similar fashion to r . However new connections are added, or old ones removed according to $[1/p]$, where $p \in (-2, 2)$ is chosen from a uniform distribution. The values on the new connections are chosen from the same initial distribution that the off diagonal values were originally set with, ie the range $\min_{i \neq j} \beta_{ij}|_{t=0}$ to $\max_{i \neq j} \beta_{ij}|_{t=0}$. Since condition (9) is computationally expensive, we use a slightly stronger criterion that is sufficient, computationally tractable yet still allows “interesting” non-definite matrix behaviour namely that the sum $\beta_{ij} + \beta_{ji}$ should be non positive.
6. μ must be positive, so should evolve according to the log-normal distribution like the diagonal components of β . Similar to β , it is a catastrophe to allow μ to become arbitrarily large. In the real world, mutation normally exists at some fixed background rate — species can reduce the level of mutation by improving their genetic repair algorithms. In ecolab, this ceiling on μ is given by the `mutation(random,maxval)` variable.

4.5 Typical Results

Figure 3 shows the time behaviour for the number of species in the ecosystem for a typical run. The phenotypic parameters were seeded randomly in the ranges $-0.005 \leq r \leq 0.01$, $-5 \times 10^{-5} \leq \beta_{\text{diag}} \leq -1 \times 10^{-4}$, $-0.001 \leq \beta_{\text{offdiag}} \leq 0.001$ and $0 \leq \mu \leq 0.09$. The r and β values were chosen so that several hundred individuals will be supported in the case of a single species system, and the offdiagonal terms large enough to permit interesting interactions between species, but not so large that the system collapsed to zero immediately. ρ was set at 10^4 , which was chosen by examining the histogram of differences between all the species. If ρ was too small, then a species’ mutant offspring would be too similar to its parent to be really a new species. This shows up as a peak at small separation values of the histogram, which shouldn’t be there according to the law of competitive exclusion.

The system rapidly evolves to one of the fixed points (by a massive extinction event!) with a negative definite β . Over time, mutations build up in the system, decreasing the stability of the system. What then follows are periods of episodic extinctions, and system growth through speciation. This is an example of *self organised criticality*[3], and gives rise to power law behaviour.

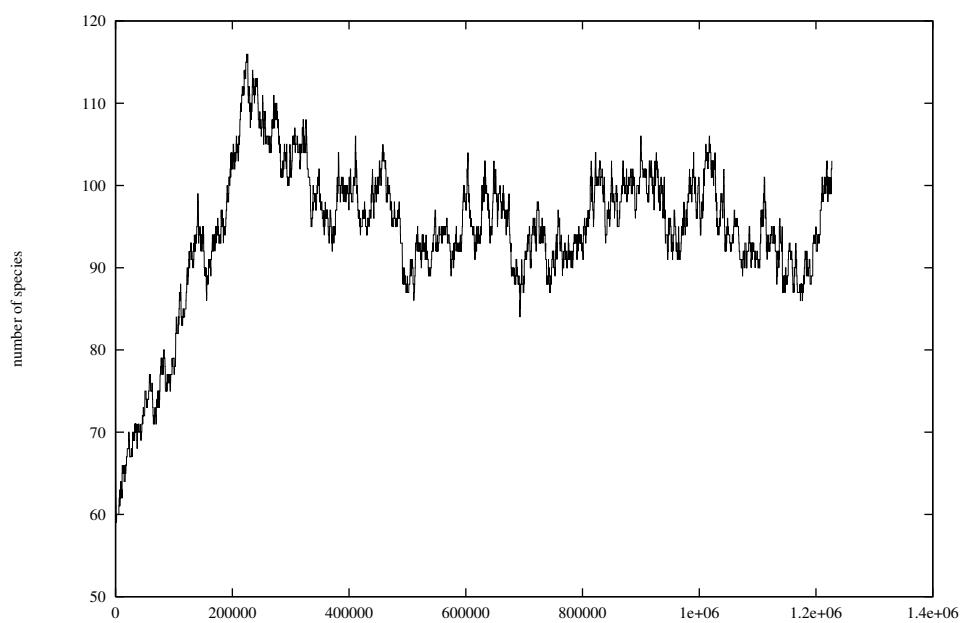


Figure 3: n_{sp} as a function of time step

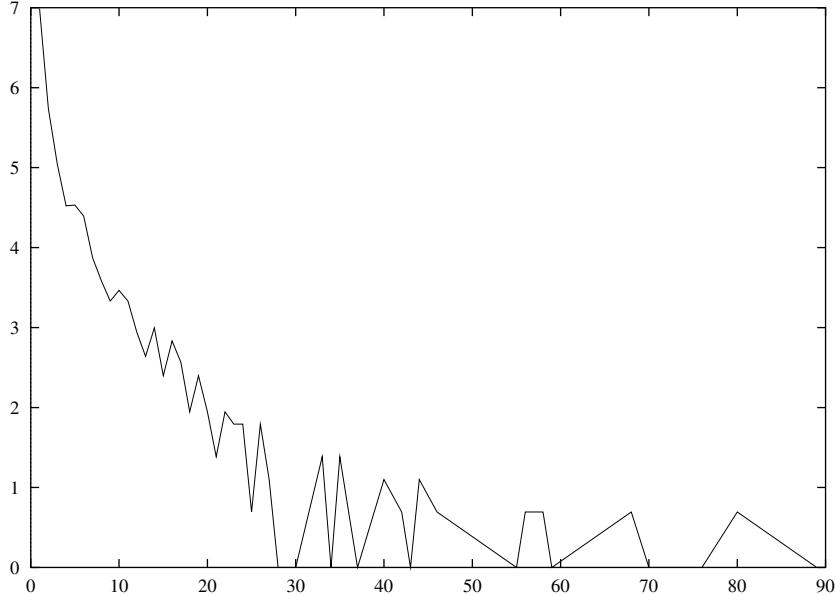


Figure 4: Distribution of species lifetimes on a log-linear plot. Distribution is unnormalised. Horizontal scale is the natural logarithm of species lifetime in timesteps.

Do we see the same power law behaviour observed by others? [4, 1] The answer is emphatically yes. If speciation and extinction events occurred uniformly throughout history, as one might naïvely expect, one would expect a Poisson distribution for species lifetimes. On a log-linear plot, this would be a straight line. Alternatively, if a power law spectrum was evident, the log-log plot would be straight. The two plots are shown in figures 4 and 5. Effectively, this is telling us that not only is there not a stable ecological equilibrium, there isn't even a steady state, whereby extinctions are balanced by speciation (a common ecological assumption).

This model then is a concrete example of the self-organised criticality predicted in these types of systems in section ???. The next section examines a possible economic model that is analogous to Ecolab, and could even be implemented using the same simulation software. It would be surprising if the dynamics weren't critically self-organised.

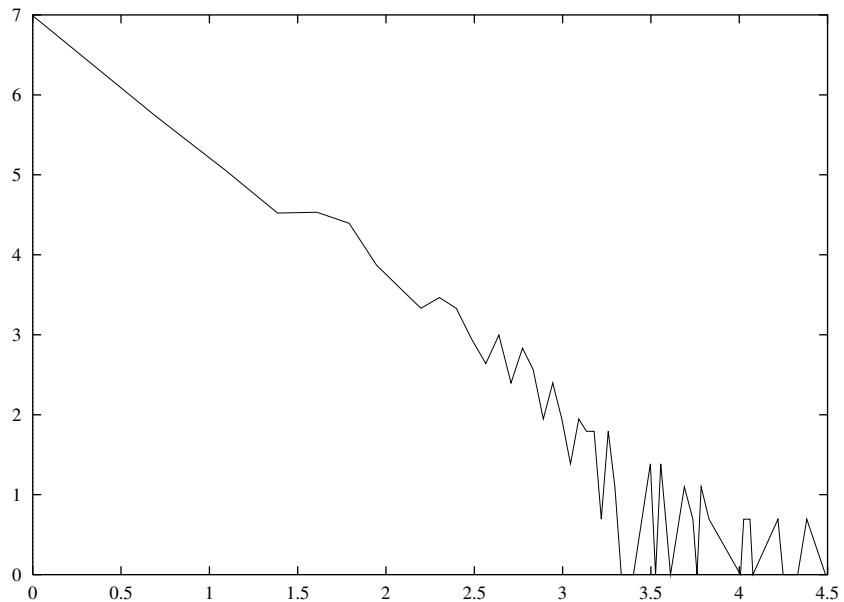


Figure 5: Distribution of species lifetimes on a log-log plot. Vertical scale is the natural logarithm of that in figure 4. Horizontal scale is the same.

5 Building an Economic Dynamics

Many inferential similarities can be drawn between the biological evolutionary model of Ecolab and the processes of a capitalist economy. The obvious analogy for a biological species is a product, and for Darwinian evolution the process of technological change. I consider a model economics (*Econolab*) based on the insights of von Neumann, one of the founders of complexity theory, who introduced *von Neumann Technology* in the late 1930s[32, 33]. In this model economy, there is a set of *commodities* labeled $i \in \mathcal{N} = \{1 \dots N\}$, and a set of *technologies* or *processes* labeled $m \in \mathcal{M} = \{1 \dots M\}$. Each process has an activity z_m , input coefficients a_{mi} and output coefficients b_{mi} , such that in one time step, $a_{mi}z_m$ of commodity i (amongst others) is consumed to produce $b_{mj}z_j$ of commodity j (amongst others). The coefficients a_{mi} and b_{mi} may be zero for some values of m and i , corresponding respectively to processes that do not require a particular input, or do not produce a particular output. This differs from von Neumann's original approach, and is more in line with that of Kemeny, Morgenstern and Thompson[14]. Blatt[5] gives a good introduction to this model, discussing its flexibility in dealing with a range of economic processes. In the words of Blatt (p67):

The von Neumann work is a great achievement of mathematical model building in dynamic economics. It is the best available theory of capital and of rate of return.

That said, there are many issues of significance in capitalism which are not captured in the von Neumann method, and which cannot be modeled in an initial rendition of Econolab. These include effective demand[?], income distribution, variable capacity and utilisation, credit and debt[?, ?].

To relate von Neumann's work back to the Ecolab ecological model, the input/output coefficients a_{mi}/b_{mi} are fixed like the r_i , β_{ij} or equation (1), and z_m is a free variable like n_i . In von Neumann's work, the dynamics is imposed in the form of an exponential growth condition:

$$z_m(t+1) = \alpha z_m(t) \quad \forall m \in \mathcal{M} \tag{12}$$

However, rather than assuming a particular form for the dynamics, we should be looking for a first order differential equation (or its difference equation equivalent) that describes the dynamics. Consider the monetary value of capital K_m associated with process m . The rate of change of this capital may be written:

$$\dot{K}_m = z_m \left(\sum_{i=1}^N b_{mi} p_i - \sum_{i=1}^N a_{mi} p_i \right), \tag{13}$$

where p_i is the price of commodity i . This has introduced two new sets of free variables K_m and p_i , for which we need to find closure relations. Clearly,

activity is limited by the availability of capital (we do not allow the possibility of credit here):

$$\sum_{i=1}^N a_{mi} p_i z_m \leq K_m \quad (14)$$

For simplicity, let us assume that each process invests a fixed proportion of its capital into production, i.e.

$$\sum_{i=1}^N a_{mi} p_i z_m = \kappa_m K_m, \quad \exists \kappa_m : 0 < \kappa_m \leq 1 \quad (15)$$

Substituting (15) into (13) gives

$$\dot{z}_m = \kappa_m \left(\frac{\sum_{i=1}^N b_{mi} p_i}{\sum_{i=1}^N a_{mi} p_i} - 1 \right) z_m. \quad (16)$$

This then, is a model dynamics analogous to the Lotka-Volterra equation (1). If price is a fixed quantity (as assumed in von Neumann theory) then (16) is equivalent to the *ansatz* (12). This is the equilibrium situation, rather like assuming that $\mathbf{n} = \hat{\mathbf{n}}$.

In reality, prices are not fixed, and must have their own dynamics. The simplest way to do this is to look for a closure relation, that relates prices to activities. The neoclassical and Austrian traditions propose that price dynamics should act as a negative feedback on the activity dynamics (eq. (16)), whereas the P-K and Sraffian tradition do not see price as an equilibrating mechanism [?]. In this work, however, we propose an ansatz on the form of the negative feedback, in a similar fashion to the ansatz used by Nosé and Hoover[10] to describe the thermostat that regulates the temperature of a non-equilibrium steady state system in a heat bath:

$$\dot{p}_i = \pi_i \left(\frac{\text{demand}}{\text{supply}} - 1 \right) p_i = \pi_i \left(\frac{\sum_{m=1}^M a_{mi} z_m}{\sum_{m=1}^M b_{mi} z_m} - 1 \right) p_i. \quad (17)$$

This differs from von Neumann, who assumes that demand never exceeds supply, and if supply exceeds demand (i.e. a surplus), then the commodity is free ($p_i = 0$). This would imply $\dot{p}_i = 0$, freezing prices. In effect this makes the system very stiff — equation (17) softens the dynamics with π_i controlling the stiffness.

6 Adding Evolution

Now that we have an economic dynamics established, we need to consider how to develop an analogy between ecological and economic evolution. By direct

analogy with Ecolab, it is clear that when a process exhausts its capital ($K_m = 0$), it forever remains that way, so this is equivalent to extinction in ecosystems. Adding new processes and commodities is conceptually easy. Blatt p57–58[5]:

What about technological progress? This can be included by assuming that the list of activities $m = 1, 2, \dots, M$ is not final, but new activities may be invented and hence become available for use, as time goes on. This makes the total number of processes a function of time: $M = M(t)$... Von Neumann himself developed his theory on the basis of an unchanged technology (all input coefficients, output coefficients and the number of processes M are constant in time), and his successors have done the same. The inclusion of technological progress appears to us to be a highly interesting avenue for further exploration.

The difficulty is deciding how to choose new coefficients a_{mi}, b_{mi}, κ_i and π_i when a new process is added. There is no genotype of a process — the closest thing to it is Dawkins's meme, and there is no genetic algorithm theory of the meme. Clearly new processes arise evolutionarily, with the new processes modeled on the old. The new coefficients will be varied randomly about the old values according to some kind of central distribution.

Recent results from Ecolab indicate that the emergent dynamics of the system is rather insensitive to the specific type of mutation algorithm chosen. Work is currently under way to classify exactly what effects different assumptions make.

In 1962, Arrow[2] pointed out that the cost per unit for production of an artifact falls as an inverse power of the number of units produced:

$$\text{cost/unit} \propto N^{-a}$$

This power law is most likely a consequence of the dynamics of technological innovation, relating to the statistical properties of the underlying “fitness” landscape, as it can be seen in Kauffman NK model[13]. Presumably an evolutionary algorithm that searches process (and commodity) space according to the same power law would be optimally matched to generating change, however another search algorithm would probably generate the same distribution of successful innovations, albeit on a different temporal scale. It should also be pointed out that large changes of process are likely to cost proportionally more than smaller changes. As any research budget is finite, the distribution of process improvements must therefore be finitely integrable (have a finite area underneath the curve), which the power law distribution is not, but the normal (Gaussian) distribution is.

7 Conclusion

Economics is clearly a dynamic process, which given its complexity will be poorly described by a linear approximation about a stable equilibrium. Rather the properties of the equilibrium will be determined by cultural evolution which operates over a longer timescale than economics. It is likely that cultural evolution will produce a self-organised critical system, and this would be one of the first questions to study. Other questions that might be looked at include looking for evidence for the Arrow law, and looking for analogues to various biological laws, such as the species-area law⁶, and dependence of biodiversity with latitude.

Perhaps the most important point I would like to make is that rather than studying a finite dimensional dynamical system, we should be studying what might be called “open dimensional” dynamical systems, where the number of degrees of freedom is finite, but not fixed at any point in time. These systems must lie between finite dimensional spaces and infinite dimensional “functional analysis” type spaces. Only then might we achieve Marshall’s economic biology, and have an understanding of why economic systems have evolved to be the way they are.

References

- [1] Chris Adami and C. Titus Brown. Evolutionary learning in the 2d artificial life system “avida”. In R. Brooks and P. Maes, editors, *Artificial Life IV*. MIT Press, 1994.
- [2] K. Arrow. The economic implications of learning by doing. *Review of Economic Studies*, 29:166, 1962.
- [3] P. Bak, C. Tang, and K. Wiesenfeld. Self-organised criticality. *Phys. Rev. A*, 38:364, 1988.
- [4] Per Bak and Kim Sneppen. Punctuated equilibrium and criticality in a simple model of evolution. *Phys. Rev. Lett.*, 71:4083, 1993.
- [5] J. M. Blatt. *Dynamic Economic Systems: A Post-Keynesian Approach*. M. E. Sharpe, New York, 1983.
- [6] E. Bruckner, W. Ebeling, J. Montano, and A. Scharnhorst. Technological innovations — a self-organisation approach. Technical Report FS II 93–302, Wissenschaftszentrum Berlin fuer Sozialforschung, Reichpietschufer 50, 1000 Berlin 30, 1993.

⁶the number of species on an island is related by a power law to the area of the island

- [7] T. J. Case. Invasion resistance, species build-up and community collapse in metapopulation models with interspecies competition. *Bio. J. Linnean Soc.*, 42:239–266, 1991.
- [8] M. W. Hirsch and S. Smale. *Differential equations, dynamical systems, and linear algebra*. Academic Press, New York, 1974.
- [9] J. Hofbauer, V. Hutson, and W. Jansen. Coexistence for systems governed by difference equations of lotka-volterra type. *J. Math. Biol.*, 25:553–570, 1987.
- [10] W. G. Hoover. *Phys. Rev. A*, 31:1695, 1985.
- [11] W. Jansen. A permanence theorem for replicator and lotka-volterra systems. *J. Math. Biol.*, 25:411–422, 1987.
- [12] S. Kauffman and W. Macready. Technological evolution and adaptative organizations. *Complexity*, 1:26, 1995.
- [13] Stuart A. Kauffman. *The Origins of Order: Self Organization and Selection in Evolution*. Oxford UP, 1993.
- [14] J. G. Kemeny, O. Morgenstern, and G. L. Thompson. A generalization of the von neumann model of an expanding economy. *Econometrica*, 24:115–135, 1956.
- [15] Richard Law and Jerry C. Blackford. Self-assembling food webs: A global viewpoint of coexistence of species in lotka-volterra communities. *Ecology*, 73:567–578, 1992.
- [16] Alfred Marshall. *Principles of Economics*. Macmillan, London, 8th edition, 1920.
- [17] R. M. May. Simple mathematical model with very complicated dynamics. *Nature*, 26:457, 1976.
- [18] Robert M. May. *Stability and Complexity in Model Ecosystems*. Princeton University Press, Princeton, New Jersey, 1974.
- [19] J. Maynard Smith. *Models in Ecology*. Cambridge University Press, London, 1974.
- [20] Richard Muller. *Nemesis: The Death Star*. Heinemann, London, 1989.
- [21] J. K. Ousterhout. *Tcl and the Tk Toolkit*. Addison-Wesley, 1994.
- [22] David M. Raup. *The Nemesis Affair*. Norton, New York, 1986.
- [23] David M. Raup. *Extinction: Bad Genes or Bad Luck?* Norton, New York, 1991.

- [24] David M. Raup and J. John Sepkowski, Jr. Periodic extinctions of families and genera. *Science*, 231:833–836, 1986.
- [25] M. Schroeder. *Fractals, Chaos, Power Laws*. Freeman, New York, 1991.
- [26] J. John Sepkowski, Jr. The taxonomic structure of periodic extinction. In Virgil L. Sharpton and Peter D. Ward, editors, *Global Catastrophes in Earth History*, number 247 in The Geological Society of America Special Papers, pages 33–44. The Geological Society of America, Boulder, Colorado, 1990.
- [27] R. K. Standish. Ecolab: Where to now? In R. Stocker, H. Jelinek, B. Durnota, and T. Bossomeier, editors, *Complex Systems: From Local Interaction to Global Phenomena*. IOS, 1996. also Proceedings of AlifeV, Nara Japan, May 1996 and *Complexity International*, vol. 3, <http://www.csu.edu.au/ci>.
- [28] Russell K. Standish. Ecolab documentation. Available at <http://parallel.acsu.unsw.edu.au/rks/ecolab.html>.
- [29] Russell K. Standish. Population models with random embryologies as a paradigm for evolution. In Russel J. Stonier and Xing Huo Yu, editors, *Complex Systems: Mechanism of Adaption*. IOS Press, Amsterdam, 1994. also *Complexity International*, vol. 2, <http://www.csu.edu.au/ci>.
- [30] C. Strobeck. n species competition. *Ecology*, 54:650–654, 1973.
- [31] C. J. Thompson. Chaos in economics and management. In R. L. Dewar and B. I Henry, editors, *Nonlinear Dynamics and Chaos*, pages 213–229. World Scientific, Singapore, 1992.
- [32] J. von Neumann. A model of general economic equilibrium and a generalization of brouwer’s fixed point theorem. In K. Menger, editor, *Ergebnisse Eines Mathematischen Kolloquiums*. Vienna, 1937.
- [33] J. von Neumann. A model of general economic equilibrium. *Review of Economic Studies*, 13:1–9, 1945.